SimDoc: Topic Sequence Alignment based Document Similarity Framework

Gaurav Maheshwari
University of Bonn, Germany
gaurav_maheshwari@outlook.com

Priyansh Trivedi
University of Bonn, Germany
pc.priyansh@gmail.com

Harshita Sahijwani
DA-IICT, India
hjsahijwani@gmail.com

Kunal Jha
University of Bonn, Germany
kunal94jha@gmail.com

Sourish Dasgupta
Rygbee Pvt. Ltd., India
sourish@rygbee.com

Jens Lehmann
University of Bonn and
Fraunhofer IAIS, Germany
lehmann@uni-bonn.de

arXiv:1611.04822v1 [cs.CL] 15 Nov 2016

ABSTRACT

Document similarity is the problem of formally representing textual documents and then proposing a similarity measure that can be used to compute the linguistic similarity between two documents. Accurate document similarity computation improves many enterprise relevant tasks such as document clustering, text mining, and question-answering. Most contemporary techniques employ bag-of-words (BoW) based document representation models. In this paper, we show that a document's thematic flow, which is often disregarded by bag-of-word techniques, is pivotal in estimating their semantic similarity. In this direction, we propose a novel semantic document similarity framework, called SimDoc. We model documents as topic-sequences, where topics represent latent generative clusters of relative words. We then use a sequence alignment algorithm, that has been adapted from the Smith-Waterman gene-sequencing algorithm, to estimate their semantic similarity. For similarity computation at a finer granularity, we tune the alignment algorithm by integrating it with a word embedding matrix based topic-to-topic similarity measure. A document level similarity score is then computed by again using the sequence alignment algorithm over all sentence pairs. In our experiments, we see that SimDoc outperforms many contemporary bag-of-words techniques in accurately computing document similarity, and on practical applications such as document clustering.

Keywords

Document Similarity, *Semantic Text Similarity*, *Topic Modeling*

1. INTRODUCTION

Document similarity measures quantify the degree of semantic similarity between a pair of documents. These measures are usually modeled as functions which map two formally represented text documents to a real space. This function modeling is based on extraction of selected textual features that are either influenced by the *hypothesis of distributional semantics* (which is primarily a keyword based statistical approach) [21] or by the *principle of compositionality* as favored in compositional semantics [15]. In recent times hybrid approaches, based on the principle of *compositional distributional semantics*, have been adopted as well [5].

An effective similarity measure has a vast variety of applications. It can power content-based recommender systems, plagiarism detection [13], and document clustering systems. Further it can be utilized for various complex NLP tasks such as paraphrase identification [4], and textual entailment [22]. Modeling document similarity measure, however, is a non-trivial research problem. This is primarily because text documents do not rigidly follow any grammatical structure or formal semantic theory. Moreover, there are several linguistic issues such as sentence paraphrasing, active/passive narration and syntactic variations, that a system needs to take in account.

Most contemporary document similarity approaches model documents as *bag-of-words* (*BoW*). While many BoW based modeling techniques (specifically, *bag-of-topics* ones) can capture the "*themes*" (represented as latent variables/vectors) of a document, they lack in representing its "*thematic flow*" (also refered to as *discourse*). Thematic flow is pivotal in computing semantic similarity between documents. A bag-of-topics approach will incorrectly result in high similarity score between two documents having similar themes occurring in a different order. This can be illustrated by the following pair of sentences, "*John loves dogs, but is scared of the cat.*" and "*The cat loves John, but is scared of dogs.*" Although both the sentences express the relationship between John and pet animals, yet they are semantically not similar. Contemporary BoW techniques would still, evaluate these two sentences to be highly similar. To this end, we propose a novel topic-modeling based document similarity framework, called *SimDoc*. We use Latent Dirichlet Allocation (LDA) [3] topic model to represent documents as sequences of topics. We also compute the semantic similarity between LDA topics using a pre-trained word embedding matrix. With the help of these two information, we calculate alignment scores between these sequences using a sequence alignment algorithm, which is an adaptation of SmithWaterman algorithm - a gene/protein sequence alignment algorithm [19]. Finally, a global semantic similarity between document pairs is computed using the sequence alignment algorithm, again over sentence tokens.

We show empirically, using data set provided by [7], that the proposed system is capable of accurately calculating the document similarity. In this experiment, SimDoc performs the second best

ACM ISBN 978-1-4503-2138-9.
DOI: 10.1145/1235

amongst every system, with an accuracy of 72.69%. We also analyse various internal components and their effect on SimDoc's overall performance.

The contributions of this paper are as follows:

- A novel document semantic similary framework, called *SimDoc*, is proposed that models document similarity as a topic-sequence alignment problem, where topic-sequence represents the latent generative representation of a document.
- A novel sequence alignment computation algorithm has been used, which is an adaptation of the popular Smith-Waterman algorithm (used for gene/protein sequence matching).
- Extensive document similarity accuracy evaluation of SimDoc, using para2vec dataset has been outlined. We also detail extensive document clustering evaluation of SimDoc using human-benchmarked document-clusters on 20Newsgroup 18828, Reuters 21578, WebKB, TREC 2006 Genome Track.

The remaining sections of the paper are organized as follows: Section 2 (*related work*) where we describe research works in textual similarity; Section 3 (*prelimaries*) that introduces the problem of document similarity, and other concepts that are used in the following sections; Section 4 (*approach*) wherein the SimDoc architecture and formulation have been described; Section 5 (*evaluation*) of SimDoc accuracy; and the Section 6 (*conclusion*) where we briefly describe the current status of SimDoc.

2. RELATED WORK

2.1 Short-text Similarity Approaches

In this section we introduce some recent works on short text similarity measures. The SemEval Semantic Textual Similarity (STS) task series has served as a gold-standard platform for computing short text similarity with a publicly available corpus, consisting of 14,000 sentence pairs developed over four years, along with human annotations of similarity for each pair [1]. In accordance to the results of SemEval 2015, team DLS@CU bagged the first position in their supervised and unsupervised run in STS [20]. The team's unsupervised system was based on *word alignment* where semantically related terms across two sentences are first aligned and later their semantic similarity is computed as a monotonically increasing function of degree of alignment. Their supervised version used cosine similarity between the vector representations of the two sentences, along with the output of the unsupervised systems. However, the underlying computation is extremely expensive and not suitable for online document similarity use cases. Another system, called Exb Themesis [9], ranked second and was the best multilingual systems amongst all the participants. The system combines vector space model [18], word alignment, and Machine Learning, implementing a complex alignment algorithm that primarily focused on named entities, temporal expressions, measurement expression, and negation handling. It tackles the problem of data sparseness and the insufficiency of overlaps between sentences through word embeddings, while integrating WordNet and ConceptNet[1] into their systems.There is significant scope for improvement for application on larger documents for all these systems. SimDoc's accuracy cannot be compared with these approaches' in short text similarity measurement tasks because SimDoc has been tailored to measure similarity in long texts.

[1]http://conceptnet5.media.mit.edu/

2.2 Long-text Similarity Approaches

Many document similarity measures are based on the seminal vector space model, where documents are represented as weighted high-dimensional vectors. This model is also popularly known as *bag-of-words* model, where words are commonly used as features. However, this model fails to capture word ordering (which adds significantly to the document semantics). It also ignores the semantics of the words. Jaccard similarity[11] which treats documents as sets of tokens, and Okapi BM25[17] and Lucene Similarity[12], which rely on term frequency and inverse document frequency of words in the documents, are other widely used document similarity measures. These measures also have the same limitations as the bag-of-words model based ones. One of the most popular techniques used for document similarity is Explicit Semantic Analysis (ESA) [8]. The rationale behind ESA is that knowledge banks, such as Wikipedia, can serve to index documents with Wikipedia articles having certain lexical overlap. ESA is particularly useful when the contextual information is insufficient (which is quite common in shorter documents). Extensions of ESA have been proposed. A prominent approach, proposed in [10], is based on measuring similarity at at both the lexical and semantic levels. Concepts are identified within documents, and then semantic relations established between concept groups (at a topic level) and concepts. It uses supervised machine learning techniques to automatically learn the document similarity measure from human judgments, using concepts and their semantic relations as features. However, ESA is still based on a sparse representation of documents and hence, may at times be quite inaccurate. Alternative document similarity techniques have been proposed that are based on statistical topic models [2, 14]. These approaches identify groups of terms (i.e. latent topics) that are strongly associated (in a distributional sense) with one another within a given document corpus. A very recent word2vec based technique, called *paragraph vector*, is proposed in [6]. It uses an unsupervised algorithm that learns fixed-length feature representations from variable-length pieces of texts, such as sentences, paragraphs, and documents. Each document corresponds to a dense vector which is trained to predict words in the document. It has been shown empirically that paragraph vectors outperform bag-of-words based measures.

3. BACKGROUND

3.1 Problem Statement

Given a pair of textual documents[2], D_1 and D_2, a document similarity measure should be able to represent the text content of the documents using some formal method and compute the semantic similarity score between the documents. Here, by semantic similarity we refer to the closeness in the their semantic content, as opposed to syntactical closeness.

Document similarity scores are usually defined on the real space $\mathbb{R}$ via the function σ defined as follows:

Definition 1 (**Document Similarity Measure** (σ)):
$\sigma : \bar{D} \times \bar{D} \mapsto [a, b]$; where,

- $\bar{D}$ is the formal representation of an input text document D.
- $a \in \mathbb{R}$ is the lower-bound score.
- $b \in \mathbb{R}$ is the upper-bound score.

[2]The generalized problem statement also includes matching documents having multimedia content; which is beyond the scope of the current paper.

Usually, these bounds are normalized to the interval [0, 1]. In the subsequent sections we first define certain foundational algorithms that form the motivation for the design of SimDoc.

3.2 Probabilistic Topic Modeling

Documents can be represented as BoW, following the assumption of *exchangeability* [3]. The assumption states that if words are modeled as Bernoulli variables, then within any random sample sequence they are conditionally independent, where the word variables are conditioned on a specific set of latent random variables called *topics*. This renders the joint distribution of every sample sequence permutation (i.e. the document variable) to remain equal, provided the topic variables are given. In other words, the assumption is that, the order of words representing a document does not matter as long as the topics, which "*generates*" the occurrence of words, are known. However, interestingly, these topics are hidden (in terms of their distributions) and hence, we need a mechanism to discover (i.e. learn) them. This learning process is called *topic-modeling*. In this paper, we use a widely adopted probabilistic topic-modeling technique, called *Latent Dirichlet Allocation*, that involves an iterative Bayesian topic assignment process via variational inferencing, over a train-corpus. The number of topics (and other related hyperparameters) needs to be preset. The prior distribution of topics over documents (and also, words over topics) is taken as Dirichlet. The process results in groupings of words that are related to each other *thematically* (in the distributional semantics sense). As an example, "*house*" and "*rent*", after the learning process, might be within the same topic.

3.3 Smith Waterman Algorithm

Smith-Waterman algorithm is widely adopted to calculate gene/ protein sequence alignment - a very important problem in the field of bio-informatics [19]. Interestingly, we can use it to quantify the degree to which two sequences of tokens, say S_1 and S_2, are aligned. It uses dynamic programming to determine the sequence-segment of S_1 that is optimally aligned with S_2 (or vice-versa). During alignment, the algorithm can either insert, delete, or substitute a token whenever a token mismatch is found during comparison. This way, one of the sequences can be transformed into the other sequence. However, editing comes with a penalty. The penalties for insertion, deletion or substitution are collectively called *gap penalty scheme*. In this paper, we have proposed a more flexible penalty scheme uses a *similarity matrix*, which can take into account the degree of similarity between the tokens. The final objective of the algorithm is to accrue a minimum penalty during editing, thereby getting the optimal sequence of edits.

Smith-Waterman algorithm differs from Levenshtein and other edit-distance based algorithms in that, it performs matching within a local "*contextual window*" (called *segment*). In the context of document similarity problem, a segment may mean a sub-sequence of words, or a sentence, or a paragraph, or a sub-sequence of sentences/paragraphs. Continuous mismatches are penalized more than ad-hoc mismatches. As an example, suppose we are interested in measuring similarity with the sentence "*John loves cats but does not love dogs*". In this case, the sentence "*John owns donkeys but does not love dogs*", will be penalized more than the sentence "*John owns cats but does not own dogs*" when they are both compared to the first sentence, because the second has two continuous mismatches, even though they both require two edits. In this way Smith-Waterman algorithm can be very useful to model the deviation in "*thematic flow*" of a discourse within a document. In section 4, we explain how LDA-based topic modeling (which generates a bag-of-topics, rather than a sequence) is integrated with the Smith-Waterman algorithm. We further adapt this algorithm to compute semantic similarity between sentences by integrating word-embeddings based similarity matrix to introduce a notion of *variable degree of token similarity*.

4. APPROACH

The SimDoc framework has five core modules: (i) *Topic-Model Learner*, (ii) *Topic-Sequence Inferencer*, (iii) *Token-level Similarity Scorer*, (iv) *Sentence-level Similarity Scorer*, and (v) *Document-level Similarity Scorer* (see Figure 1).

The Topic-Model Learner module receives a training document corpus, and encodes each training document into a n-dimensional vector $[p_1,\ p_2,\ p_3\ ...p_n]$. Here n represents the number of topics in the trained LDA model and each value p_i represents the probability of the document to have i^{th} topic. After the trained model is generated, it is then used by the Topic-Sequence Inferencer to represent a given document as a *sequence* of latent topics using the vector representation of documents and a word-to-topic inverted index.

To compute similarity between two documents, their corresponding topic-sequence representations are fed into the Sentence-level Similarity Scorer, which uses an adaptation of Smith-Waterman alignment algorithm (discussed in section 3.3). For every mismatch during alignment computation, the algorithm uses Token-level Similarity scorer as a novel compensation computing module that helps to evaluate the degree of mismatch between two topics. It uses cosine similarity between topic-to-vec representation of every topic, where topic-to-vec is the average word-embedding (i.e. a high-dimensional vector represention [16]) of top-k most probable words in that topic.

After similarity is computed between a pair of topic sub-sequence (i.e. pair of sentences from the two documents), the Document-level Similarity Scorer finally computes the similarity at the document level, by applying the same proposed alignment algorithm (as explained in Sec 3.3) by representing documents as *topic sequence segments*, where every segment represents its corresponding sentence. The compensation during mismatch of two topic-sequence-segments is calculated by the pre-computed sentence-to-sentence similarity done by Sentence-level Similarity Scorer.

4.1 Topic-Model Learner

This is a training-phase module that learns topic-distributions from each document (and thereby learns the word-distribution for each topic) in the train-corpus. We use Latent Dirichlet Allocation (LDA) (Section 3.2) based topic modeling for our purpose. It is to be noted that an LDA-based topic model is more accurate when trained over a fixed domain that has a particular vocabulary pattern (i.e. domain-specific linguistic variations and jargon). For instance, a topic model trained over documents from the area of computer science cannot be used to accurately generate topic distributions of documents containing travel blogs. However, it might be able to perform relatively better in related fields such as electrical engineering or statistics or mathematics.

The Topic-Model Learner first performs text pre-processing on the train-corpus which includes tokenization, lemmatization, and stop-word removal. This pre-processing ensures that the LDA model is trained over a condensed natural language text devoid of words which add little or no semantic value to any document. All the pre-processing tasks are done using Spacy[3]. The train-corpus documents are then passed through Gensim's[4] implementation of LDA to learn topic-distributions for the documents.

[3]https://spacy.io/

[4]https://radimrehurek.com/gensim/

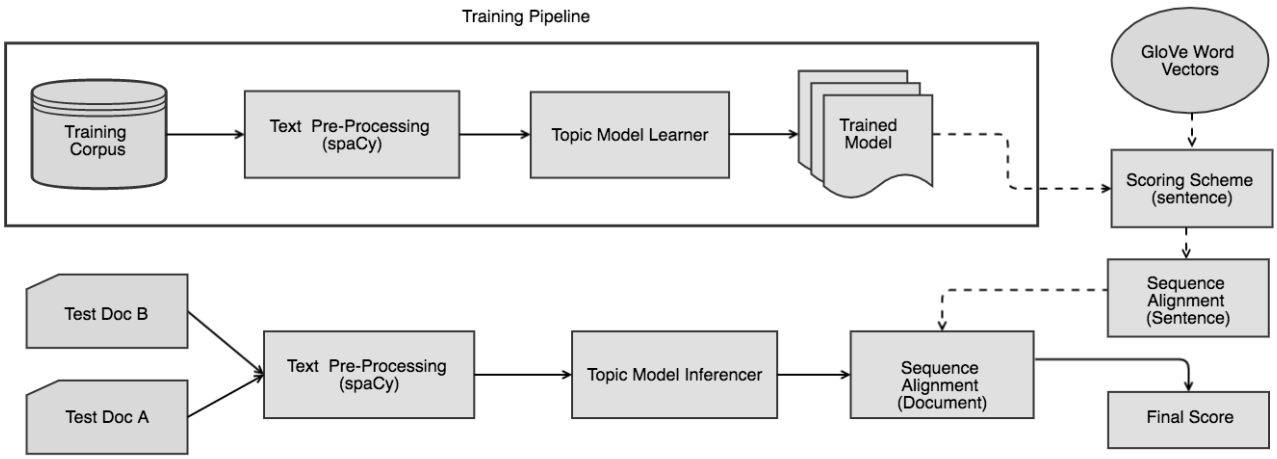

Figure 1: SimDoc Architecture

This module also creates an inverted topic-word distribution index that maps each word of the vocabulary to topics, along with the probability of that word in the corresponding topic. Its utility is explained in the section below.

4.2 Topic-Sequence Inferencer

This is an inferencing phase module. When each document of an unseen document pair is fed into the module, it first performs the same NLP pre-processing as the Topic-Model Learner module. After that, it performs voice normalization on every sentence in the documents, thereby converting passive sentences into their active form. Without this normalization step, the thematic flow of similar sentences (and hence, documents) will appear different even if they have the same semantic content.

The cleaned document pair is fed into Gensim's trained LDA topic model to infer topic distributions of the documents. Thereafter, the module transforms the documents into their *topic sequence* based representations. The word-to-topic mapping is done by using the inverted topic-word distribution index (described in the previous section) where, as the document is passed through the model, every word in the document is assigned the maximum probable topic. The generated topic-sequence represents the transition from one semantic theme to the other (i.e. the "*thematic flow*" of the document content).

Further, the module divides a topic sequence into "*topic sequence segments*", where a segment represents a sentence. Sentence segmentation is important because of two reasons. First, to capture the discourse-level locality of a semantic similarity match[5], it is important to consider sentence-boundary based topic-sequence-segments (rather than longer topic-sequences).

Secondly, in long topic-sequences without sentence segmentation, early penalty due to sentence mismatches propagates cumulatively, thereby adversely affecting later stage sentence matches.

4.3 Token-level Similarity Scorer

This module is responsible for computing compensation whenever a topic-to-topic mismatch occurs while computing the alignment score between two topic-sequence-segments (i.e. sentences). The resultant score is expected to represent the degree of closeness between two topics (based on their constituent top-k words). For example, if the top-4 words of three topics t_1, t_2, t_3 are ["lion", "cub", "flesh", "wild"], ["insect", "ants", "forest", "ferns"] and ["kindergarten", "toddler", "alphabets", "cubs"]; the score for the t_1, t_2 pair should be higher than that of the t_1, t_3 pair.

To accomplish this, we encode each topic into a vector space by first transforming its corresponding top-k words into high-dimensional vectors (i.e. word embeddings) using GloVe based pre-trained word vectors [16]. This model is trained over a specific domain corpus that best suits the document pairs. In case the domain of the documents is unknown, we may train the model over a generic corpus such as Wikipedia. The topic vector is then computed as an average of the top-k word vectors. The semantic similarity between two topics can then be computed by calculating the cosine similarity between them.

For every topic t_i in the trained LDA model, let w_{i1}, w_{i2}, ... w_{in} be its top-n words, and K be the number of total topics in the model. Let G be the vector space representation of the GloVe matrix. Let $Enc_word : (w_{ij}, G) \rightarrow wv_{ij}$, where $wv_{ij} \in G$ be an encoding function mapping word tokens to their vector representations. Using that, we define:
$Enc_topic(t_i) = \frac{1}{n} \sum_{j=0}^{n} \big(Enc_word(w_{ij}, G)\big)$ to be the vector encoding function for the i^{th} topic, where $i \in [0, K]$.
Then, the function responsible for computing topic-to-topic similarity can be defined as:

$$topic_similarity(t_i, t_j) = \frac{Enc_topic(t_i) \cdot Enc_topic(t_j)}{||Enc_topic(t_i)||\ ||Enc_topic(t_j)||}$$

where $i, j \in [0, K]$

4.4 Sentence-level Similarity Scorer

This module computes the similarity between a pair of topic-sequence-segments (a topic-sequence-segment represents one sentence of a document, as discussed in Section 4.2). We use an adaptation of the Smith Waterman algorithm, which in turn uses the Token Level similarity scorer, in order to compute the alignment score (or conversely, the degree of disalignment) between the topic-sequence-segments. Before formalizing the algorithm, we first describe some preliminary concepts as follows:

- S_i^a is the i^{th} topic sequence segment (correspondingly, representing the i^{th} sentence) of document D_a.
- $token_x^{S_i^a}$ is the token on x^{th} position in S_i^a. $token_x^{S_i^a} \in$

[5]See Section 3.3 for an example illustrating the effect of localized alignment computation within a sentence

$[t_1, t_2, t_3..t_n]$ (Topics in the LDA Model).

- $sentence_similarity(S_i^a, S_j^b)$ is a function which computes the semantic similarity between i^{th} topic-sequence segment of document D_a and j^{th} topic sequence segment of document D_b. $sentence_similarity : (S_i^a, S_j^b) \rightarrow [0, 1]$; where 1 is the maximum possible similarity between two segments.

- $score(token_x^{S_i^a}, token_y^{S_j^b})$ is the score assigned by the sequence alignment algorithm when comparing two tokens of the sequence. As discussed in Sec. 3.3, there can either be a match or a mismatch between the tokens. Every match accrues a reward. For a mismatch, there can be three types of edit possible - insertion, deletion, and substitution. Each of these edits comes with a penalty (i.e. cost of edit). A scoring scheme is responsible for deciding these penalties. This algorithm uses a linear combination of the edit penalty (called Gap Penalty), and the topic pair's similarity computed by the Topic-level Similarity Scorer. The scoring scheme is defined as follows:
  $score(token_x^{S_i^a}, token_y^{S_j^b}, op) = G_{op} + (f \times topic_similarity(t_i, t_j))$
  where:
  - $f \in [0, 1]$ is a discount factor for the similarity score.
  - $op \in [Ins, Sub, Del]$ signifies the edit operation for which this score is to be computed.
  - $G_{op} \in \mathbb{R}$ (negative real numbers) is the Gap Penalty for an edit, and thus can take three different values, represented by $G_{ins}, G_{sub}, G_{del}$.

- $value(token_x^{S_i^a}, token_y^{S_j^b})$ is the *cumulative* alignment score assigned to the topic sequence segments till x^{th} token in S_i^a sequence and y^{th} token in S_j^b sequence. It is described below (as a part of algorithm description) below.

For better readability, we will hereupon refer to the function $score(token_x^{S_i^a}, token_y^{S_j^b}, op)$ as $s(x, y, op)$,

and $value(token_x^{S_i^a}, token_y^{S_j^b})$ as $v(x, y)$. We define our proposed sequence alignment algorithm by the following Bellman equations:

$$v(x,y) = \begin{cases} 0 & \text{iff x = 0 or y = 0} \\ max\big(0,\ v(x-1, y-1) + M\big) & \text{iff } token_x^{S_i^a} = token_y^{S_j^b} \\ max \begin{cases} 0 \\ v(x-1, y) + s(x, y, Del) & \text{iff } token_x^{S_i^a} \neq token_y^{S_j^b} \\ v(x, y-1) + s(x, y, Ins) & \text{iff } token_x^{S_i^a} \neq token_y^{S_j^b} \\ v(x-1, y-1) + s(x, y, Sub) & \text{iff } token_x^{S_i^a} \neq token_y^{S_j^b} \end{cases} \end{cases}$$

for $x \in [0, m]$ and $y \in [0, n]$; where:
m and n are lengths of S_i^a and S_j^b respectively.
M is the Match Gain (i.e. reward for a match)
$sentence_similarity((S_i^a, S_j^b) = v(m, n) / \sum_{i=1}^{max(m,n)} i$ The parameter values are empirically determined during cross-fold testing.

4.5 Document-level Similarity Scorer

Thematic discourses are often spread across more than one sentence in a document. In document pairs with high similarity, we expect to find some alignment in this discourse. To model this, we apply the same proposed sequence alignment algorithm but now, over a sequence of topic-sequence-segments (representing the order in which sentences appear in the given document). During the alignment process, the Document-level Similarity Scorer uses Sentence-level Similarity Scorer to compute the degree of mismatch between two sentences, when it is found. The algorithm can be expressed in a similar fashion as follows:

- D_a is the topic sequence representative of the document's text, which is divided into segments (sentences). $D_a = \{S_a^i\};\ i \in [0, M]$, where M is the number of sentences in the document D_a.

- S_i^a, as defined in Sec 4.4, are used as the tokens for this sequence alignment algorithm. It represents a topic-sequence-segment (sentence) of the document D_a.

- $document_similarity(D_a, D_b)$ is a function that computes the semantic similarity between two documents, D_a and D_b.

- $score_{doc}(S_i^a, S_i^b, op)$ is the score assigned by the sequence alignment algorithm when comparing two tokens of the sequence (two sentences). We use a similar scoring scheme as in Sec. 4.4, having both constant gap penalties as well as a similarity matrix.
  $score_{doc}(S_i^a, S_j^b, op) = G_{op} + (f \times sentence_similarity(S_i^a, S_j^b))$
  where:
  - f is the discount factor for the similarity score
  - op is the edit operation
  - G_{op} is the constant Gap Penalty
  - $sentence_similarity$ function is responsible for the similarity matrix based score (defined in previous section)

- $value_{doc}(S_i^a, S_j^b)$ is the *cumulative* alignment score assigned to the document pair $(D_a^{'}, D_b^{'})$ where $D_a^{'}$ is D_a counted till the i^{th} sentence, and $D_b^{'}$ is D_b counted till j^{th} sentence.

For better readability, we refer to $score_{doc}(S_i^a, S_j^b, op)$ as $s_d(i, j, op)$ and $value_{doc}(S_i^a, S_j^b)$ as $v_d(i, j)$. The Bellman equations for this algorithm are:

$$v_d(i,j) = \begin{cases} 0 & \text{iff x = 0 or y = 0} \\ max\big(0,\ v_d(i-1, j-1) + M\big) & \text{iff } S_i^a = S_j^b \\ max \begin{cases} 0 \\ v_d(i-1, j) + s_d(i, j, Del) & \text{iff } S_i^a \neq S_j^b \\ v_d(i, j-1) + s_d(i, j, Ins) & \text{iff } S_i^a \neq S_j^b \\ v_d(i-1, j-1) + s_d(i, j, Sub) & \text{iff } S_i^a \neq S_j^b \end{cases} \end{cases}$$

for $x \in [0, M] \& y \in [0, N]$, where M and N are lengths of D_a and D_b respectively.

The final document similarity is calculated as follows:
$document_similarity(D_a, D_b) = v_d(M, N) / \sum_{i=1}^{max(M,N)} i$

5. EVALUATION

We perform three experiments to measure and analyse the performance of SimDoc. The first two are text analytics based tasks: *document similarity* and *document clustering*. These tasks are meant to quantitatively evaluate the accuracy of SimDoc. The third task is meant to compare different aspects of SimDoc to better understand the effect of different modules on the overall performance of the system.

5.1 Document Similarity

5.1.1 Evaluation Setup and Dataset

In this task, the system is given a set of three documents, and is expected to detect the most similar pair of documents in the set. A dataset of 20,000 such triples were generated and made publicly available by [7]. They collected the URLs of research papers archived on arXiv.org [6]; and based on their keywords and categories, made these triples of URLs. This was done in such a way that Document 2 (D_2) has some subjects common with Document 1 (D_1) but none with Document 3 (D_3). And consequentially, the system should report that D_1, D_2 are more similar than D_2, D_3. Here, by documents, we refer to the research papers found on the URLs present in the dataset.

In our experiment, we do not fetch the entire paper from the URLs, but only their abstracts. This is done because comparing papers would require a robust text extraction module, which can detect and translate tables, equations, figures etc, which lies beyond the scope of our system. Comparing based on abstracts is a more difficult task since abstracts have relatively less structural variations, less semantic information, and are of a significantly shorter size w.r.t. entire papers. Thus, the resultant accuracy of our system is expected to improve if we compare the entire research papers instead.

Training: We train a LDA model on a large subset (1.1×10^6) of abstracts available on arXiv.org. As discussed in section 4.5, our sequence alignment algorithms have a total of 10 parameters, and the LDA topic model has three parameters: number of topics (set to 100), training iterations (25000), number of passes during every iteration (set to 6), and two hyper-parameters influencing the Dirichlet Prior: α (set to 0.1) and β (set to 0.001).
The sequence alignment parameters were decided automatically, by running a Gradient Descent over a small subset (1000 document triples), with the objective function being the percentage accuracy of the system. Thus, the 20,000 triples are divided into - *training set* comprising of first 1000 triples; and *test set* comprising of rest 19,000 triples.

5.1.2 Evaluation Goal

We measure the ability of the system to correctly deduce, with respect to human judgement, that D_2 is more similar to D_1 when compared to D_3. In other words, $document_similarity(D_2, D_1) > document_similarity(D_2, D_3)$. After the training was done we ran SimDoc over the test dataset and recorded the percentage of correctly predicted triples. We compare the performance of our system with respect to other contemporary techniques such as Jaccard Index [11], Lucene Index [12], Okapi BM25 [17], and with Para2Vec - the system proposed in [6] (the creators of this dataset).

5.1.3 Results and Analysis

As depicted in Table 1, SimDoc is able to outperform the other systems, except for [6]. This validates our hypothesis that comparing thematic flow of documents improves the accuracy of document similarity based tasks. Basic BoW based techniques fall short in this task because in many partially similar factual documents there is a considerable overlap in the vocabulary, and often in the word frequency distributions. The ability to compare the sequence of topics enables the system to get accurate results even in these cases. This is not to suggest that we achieved an upper bound on the performance of our system on the task. There are some limitations of the system which we discuss in Sec. 6.

Table 1: Document similarity task

System	Accuracy
SimDoc	72.69%
Document embedding with paragraph vectors	85% [7]
Okapi BM25 over Bag Of Words	59.08%
Lucene index over Bag Of Word	59.86%
Jaccard index over Bag Of Word	66.02%

In this experiment, Para2Vec achieves a significantly better accuracy than every other system. However, it should be taken into account that they were calculating similarity over entire research papers, while other systems were comparing only abstracts. In comparison with Research papers, which usually provide detailed information about their own hypothesis and approaches, abstracts have significantly less semantic content, and typically follow a similar structure. We observed abstract comparisons tend to give false positive for papers in same field. Thus, Computing the semantic similarity of abstracts is a much more challenging task.

5.2 Document Clustering

5.2.1 Evaluation Setup and Dataset

Another way to measure SimDoc accuracy is to observe how well it performs with respect to human clustered document datasets. We evaluate SimDoc on four different benchmark corpus datasets which we describe below:

20-Newsgroup[8]: This dataset is derived from the CMU Text Learning Group Data Archive. It involves newsgroups with a collection of 20,000 messages, collected from 20 different Internet-news groups. The dataset includes 1000 messages from each of the twenty newsgroup, chosen at random and were partitioned on the basis of group name.

Reuters 21578[9]: Widely used test set for text categorization evaluations, these documents are Reuters newswire stories, across five different content themes The five category sets are Exchanges, Orgs, People, Places and Topics. The first four above-mentioned categories correspond to named entities of the specified types and Topics categories are economic subject categories.

WebKB[10]: This dataset is derived from the World Wide Knowledge Base project of CMU text learning group. These webpages were collected from various computer science universities and manually classified into seven different classes: student, faculty, staff, department, course, project, and other.

TREC 2006 Genomics Track[11]: This dataset is derived from 49 journals for Genomics track. Journal articles range from topics of

[6] https://arxiv.org/

[7] On entire research paper and not just abstract

[8] http://qwone.com/ jason/20Newsgroups/

[9] http://www.daviddlewis.com/resources/testcollections/reuters21578/

[10] http://www.cs.cmu.edu/ webkb/

[11] http://trec.nist.gov/data/t2006_genomics.html

epidemiology, alcoholism, blood, biological chemistry to rheumatology and toxicological sciences.

These documents are clustered based on semantic similarity. In other words, two documents belonging to the same cluster are more similar, when compared to one document taken from one cluster and one from another.

5.2.2 *Evaluation Goal*

In this experiment, we try to cluster these documents based on our Document similarity measure. Let there be m clusters in the dataset and let the cluster size of each cluster be n. For each document in dataset, we use SimDoc to select top $(n-1)$ most similar documents across the entire dataset. Ideally, the top $(n-1)$ documents selected, should belong to the same cluster as the document which is being used for comparision. For each set of retrieved documents, given a particular document, we then compute the *average accuracy*, in terms of the four measures (i.e. *Precision*, *Recall*, *F-score*, and *Rejection*). Then we compute the *mean average accuracy*, in terms of each measure, over all the clusters for a given benchmark dataset.

5.2.3 *Results and Analysis*

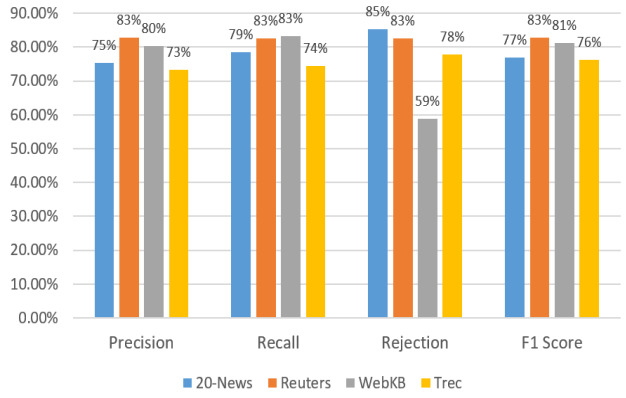

Figure 2:
Human-benchmarked Document-cluster based Accuracy Evaluation

As shown in the Figure 2, we obtain high Mean Average Precision (MAP) (20NewsGroup: 75.3%; Reuters: 82.8%; WebKB: 80.34%) and high Mean Average Recall (MAR) (20NewsGroup: 78.5%; Reuters: 82.80%; WebKB: 83.14%) for SimDoc with respect to *W-Rel*. Also, we observed that *SimDoc* works very accurately in giving a low similarity score in case of dissimilar document pair across all datasets. This has been clearly shown by the Mean Average Rejection (MA-Rejection) (20NewsGroup: 85.2%; Reuters: 82.5%; WebKB: 58.8%). However, we found certain anomalies in the behavior of SimDoc as per W-Rel. As an example, we observed that in a corpus which contained one document pertaining to *contraceptives* and another to *environment*, SimDoc incorrectly showed high semantic similarity. This is due to the co-occurrence of the terms within very similar contexts (in the above example, the observed context, *sex* and *population* respectively, are strongly mutually related). However, in such cases, SimDoc, as per S-Sim, correctly showed high value of dissimilarity, and as per S-Rel, correctly showed high relatedness between the two documents.

Table 2: Document Level Similarity Measure

Algorithm Used	Accuracy
SmithWaterman	72.69%
Root Mean Square Deviation	69.56%
Mean	66.93%

5.3 Extended Analysis

Using the experimental setup of Section 5.1 we try and evaluate different aspects of SimDoc.

I. Effect of different word embeddings: As discussed in Sec. 4.3, the similarity between two topics are computed based on a word embedding matrix. Upon replacing the pre-trained GloVe (trained by [16] on Wikipedia + Gigaword corpus, 300 dimensions) with pre-trained word2vec (trained by Google on Google News dataset, 300 dimensions) we observed that the overall performance of our system goes down from 72.69%to 69.73%. We hypothesize that apart from the way they're trained, a major reason for this change is the fact that GloVe's training set is closer to the task's dataset. Domain specific embeddings training might increase the accuracy further.

II. Effect of different Document-level Similarity Scorer: Given all-pair sentence similarity, the document level similarity can be calculated in various ways, and not just via sequence alignment algorithms. Average of all-pair similarity can be one measure, another could be to find best match in the document and compute a Root Mean Square Distance (RMSD) of these best match sentence pairs. Both these mechanism were implemented and compared with Sequence Alignment based Document Level Scorer. The average based technique performed the worst, as even in similar documents, every sentence pair doesn't have a high similarity value. Usually the underlying similarity is between the arguments and discourses. Hence, performing this average would take into account un-necessary sentence pairs. This also motivates the use of RMSD over best-matching sentence pairs. However, in our results, Document Level Sequence Alignment outperforms both of them. That's because these discourses often span across more than one sentence, and best matching pairs cannot take this into account. The results of this experiment, as shown in 2, reflects the same.

II. Effect of Topic Modeling: We aim to evaluate whether or not using Topics improves the performance of SimDoc, and to what extent. To measure that, we used an alternate implementation of the system which doesn't model the document as segmented topic sequences, but only as word vectors (encoded using GloVe). The token level similarity was computed using a simple cosine of these vectors. We observed that without word embeddings, SimDoc reaches a 64.3% accuracy. This proves that word embeddings alone are insufficient to represent the thematic flow of the documents, and topic modeling and sequencing is pivotal for this system.

III. Effect of optimizing alignment algorithm parameters: We evaluate the performance of SimDoc with sub-optimal parameters to understand their impact on the performance of the system. Taking the initial parameters of (compensation factor, match gain, insert penalty, delete penalty, substitute penalty) as (1,1,-0.5,-0.5,-1) respectivelty, for both the sequence alignment algorithm, results in 60.2% accuracy. So, by optimizing the parameters, the system's accuracy improved by almost 12%. It should be noted that in some cases, the gradient descent algorithm encountered a sub-optimal local minima, and thus it is strongly advised to use a Stochastic Gradient Descent while optimizing these parameters.

6. CONCLUSION

In this paper, we propose SimDoc- a topic sequence alignment based document similarity measure. We compared SimDoc with contemporary document similarity measures such as Jaccard, Lucene Index, BM25 and [6]. SimDoc achieved a high accuracy making it a promising paradigm of comparing documents based on their semantic content.

We can further improve SimDoc by incorporating negation-handling, dependency-parsing based complex voice normalization, named entity recognition, word sense disambiguation, and a sentence-simplification module for paraphrase normalization.

7. REFERENCES

[1] E. Agirre, C. Banea, et al. Semeval-2015 task 2: Semantic textual similarity, english, s-panish and pilot on interpretability. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015), June*, 2015.

[2] D. Blei and J. Lafferty. Correlated topic models. *Advances in neural information processing systems*, 18:147, 2006.

[3] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003.

[4] C. Brockett and W. B. Dolan. Support vector machines for paraphrase identification and corpus construction. In *Proceedings of the 3rd International Workshop on Paraphrasing*, pages 1–8, 2005.

[5] S. Clark, B. Coecke, and M. Sadrzadeh. A compositional distributional model of meaning. In *Proceedings of the Second Quantum Interaction Symposium (QI-2008)*, pages 133–140, 2008.

[6] A. M. Dai, C. Olah, and Q. V. Le. Document embedding with paragraph vectors. *arXiv preprint arXiv:1507.07998*, 2015.

[7] A. M. Dai, C. Olah, Q. V. Le, and G. S. Corrado. Document embedding with paragraph vectors. In *NIPS Deep Learning Workshop*, 2014.

[8] E. Gabrilovich and S. Markovitch. Computing semantic relatedness using wikipedia-based explicit semantic analysis. In *IJCAI*, volume 7, pages 1606–1611, 2007.

[9] C. Hänig, R. Remus, and X. De La Puente. Exb themis: Extensive feature extraction from word alignments for semantic textual similarity.

[10] L. Huang, D. Milne, E. Frank, and I. H. Witten. Learning a concept-based document similarity measure. *Journal of the American Society for Information Science and Technology*, 63(8):1593–1608, 2012.

[11] P. Jaccard. *Etude comparative de la distribution florale dans une portion des Alpes et du Jura*. Impr. Corbaz, 1901.

[12] A. Jakarta. Apache lucene-a high-performance, full-featured text search engine library, 2004.

[13] C.-H. Leung and Y.-Y. Chan. A natural language processing approach to automatic plagiarism detection. In *Proceedings of the 8th ACM SIGITE conference on Information technology education*, pages 213–218. ACM, 2007.

[14] W. Li and A. McCallum. Pachinko allocation: Dag-structured mixture models of topic correlations. In *Proceedings of the 23rd international conference on Machine learning*, pages 577–584. ACM, 2006.

[15] P. Liang, M. I. Jordan, and D. Klein. Learning dependency-based compositional semantics. *Computational Linguistics*, 39(2):389–446, 2013.

[16] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–43, 2014.

[17] S. E. Robertson, S. Walker, S. Jones, M. M. Hancock-Beaulieu, M. Gatford, et al. Okapi at trec-3. *NIST SPECIAL PUBLICATION SP*, 109:109, 1995.

[18] G. Salton, A. Wong, and C.-S. Yang. A vector space model for automatic indexing. *Communications of the ACM*, 18(11):613–620, 1975.

[19] T. F. Smith and M. S. Waterman. Identification of common molecular subsequences. *Journal of molecular biology*, 147(1):195–197, 1981.

[20] M. A. Sultan, S. Bethard, and T. Sumner. Dls@ cu: Sentence similarity from word alignment. *SemEval 2014*, page 241, 2014.

[21] P. D. Turney, P. Pantel, et al. From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research*, 37(1):141–188, 2010.

[22] N. Zeichner, J. Berant, and I. Dagan. Crowdsourcing inference-rule evaluation. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2*, pages 156–160. Association for Computational Linguistics, 2012.